

Using tensor hypercontraction density fitting to achieve an $O(L^4)$ CISD algorithm

Neil Shenvi¹, Helen Van Aggelen^{1,2}, and Weitao Yang¹

¹ *Dept. of Chemistry, Duke University, Durham, NC 27708 and*

² *Dept. of Inorganic and Physical Chemistry, Ghent University, 9000 Ghent, Belgium*

(Dated: September 14, 2012)

Recently, Hohenstein et al[1] introduced tensor hypercontraction density fitting to decompose the rank-4 electron repulsion integral tensor as the product of five rank-2 tensors. In this paper, we use this methodology to construct an algorithm which calculates the approximate ground state energy in $O(L^4)$ operations. We test our method using several small molecules and show that we quickly approach the CISD limit with a small number of auxiliary functions.

The problem of the rapid growth of the electronic wavefunction with system size has plagued quantum chemistry for decades[2, 3]. There have been many attempts to conquer the ‘curse of dimensionality’ and a large number of highly successful approximations have been developed. The difficulty inherent in electronic structure calculations is that correlated methods usually scale as high orders of the number of basis functions involved in the calculation. A straightforward implementation of Hartree-Fock scales as $O(L^4)$ where L is the number of basis functions. Approximate methods employed to capture correlation usually face a trade-off between accuracy and efficiency[4, 5]. Perturbative treatments such as MP2 and MP4 scale as $O(L^5)$ and $O(L^6)$, respectively[4]. Configuration interaction methods, which add in single-, double-, triple- and quadruple order excitations, also form a hierarchy of methods which scale as $O(L^5)$ and higher[6]. A plethora of other, highly accurate methods based on coupled-cluster theory[6], 2-particle reduced density matrices[7, 8], and reduced active space diagonalization[9] also scale as at least $O(L^6)$. Due to this high scaling, these methods are generally not applicable beyond small molecules. Instead, quantum chemists have increasingly turned to density functional theory (DFT), which offers low $O(L^4)$ or even $O(L^3)$ formal scaling while managing to capture varying amounts of the electronic correlation[10]. Nonetheless, the search for efficient wavefunction-based methods that are competitive with DFT has continued, motivated by the importance of strong correlation in many systems such as transition metal clusters, solid state devices, and molecules far from their equilibrium geometries.

One approach which attempts to reduce both the cost and scaling of correlated electronic structure methods is the decomposition of the electronic repulsion integral (ERI) tensor, which is naturally a rank-4 object, into products of lower-rank objects. Resolution-of-the-identity techniques[11–14] are one subset of this approach, as are pseudospectral methods[15–17]. Our algorithm was strongly motivated by the recent work of Hohenstein et al.[1], who developed a method which they named ‘tensor hypercontraction density fitting’ (THC-DF). The authors showed that their decomposition could be used to derive $O(L^4)$ MP2 and MP3 algorithms and suggested that it could also be used to improve the efficiency of a wide variety of electronic structure methods. Here, we apply their idea to the CISD method and show that it yields a dramatic reduction of computational cost, from $O(L^6)$ for normal CISD to $O(L^4)$ for our method. Our method will rigorously approach the traditional CISD energy as the number of auxiliary functions is increased. Therefore, our method will face the same challenges, such as lack of size-extensivity, as the traditional CISD method. However, we believe that the substantial reduction in computational scaling exhibited by our method more than makes up for the well-known shortcomings of CISD, particularly given the fact that numerous methods exist for the correction of these shortcomings[18].

The electronic Hamiltonian of an atom or molecule can be written in 2nd-quantized notation as

$$\hat{H} = \hat{H}_1 + \hat{H}_2 \quad (1)$$

$$= \sum_{ik} h_k^i \hat{c}_i^\dagger \hat{c}_k + \sum_{ijkl} \epsilon_{jl}^{ik} \hat{c}_i^\dagger \hat{c}_k \hat{c}_j^\dagger \hat{c}_l \quad (2)$$

where h_k^i is the rank-2 matrix of 1-electron integrals and $\epsilon_{kl}^{ij} = (ik|jl)$ is the rank-4 electronic repulsion integral tensor. The ERI can be decomposed using a set of P_H auxiliary functions as

$$\epsilon_{jl}^{ik} = \sum_{a,b=1}^{P_H} x_{ia} x_{ka} Z_{ab} x_{jb} x_{lb}. \quad (3)$$

This decomposition can be achieved trivially by writing ϵ as a $L^2 \times L^2$ matrix and diagonalizing it to obtain eigenvector v_{ik}^λ and eigenvalues d^λ . We can then pick some arbitrary orthonormal 1-electron basis x_{ia} and decompose the

eigenvectors v_{ik}^λ into a product of rank-1 objects by solving L^2 linear equations

$$v_{ik}^\lambda = \sum_{a=1}^{L^2} y_a^\lambda x_{ia} x_{ka}. \quad (4)$$

Letting

$$Z_{ab} = \sum_{\lambda=1}^{L^2} y_a^\lambda d^\lambda y_b^\lambda \quad (5)$$

we obtain the exact decomposition in Eq. (3). The problem with this decomposition is that 1) it involves a large number of auxiliary basis functions $P_H = L^2$ and 2) it is based on the diagonalization of the $L^2 \times L^2$ matrix ϵ , which requires $O(L^6)$ operations.

One of the most important conclusions of [1] was that real molecular ERI tensors can be well-approximated to mH accuracy by the expression in Eq. (3) with a number of auxiliary functions that scale only as $P_H = O(L)$. The authors also devised an algorithm to obtain the decomposition in Eq. (3) in $O(L^4)$ operations. Although this previous result is crucial to ensuring the low scaling of our overall algorithm, for the purposes of testing the intrinsic accuracy of our method, we will use the exact decomposition prescribed in Eq. (5) to avoid any errors inherent to the approximation of the ERI tensor. Another important feature of the THC-DF decomposition is that it enables the efficient $O(L^3)$ transformation of the ERI from the atomic orbital basis to any other arbitrary one-electron basis, such as the orthonormal molecular orbital basis. Because of this simple transformation property, we will assume that the ERI is already written in an orthonormal basis such as the molecular orbital basis.

Our CISD method itself is based on a parametrization of a single-Slater determinant reference wavefunction, usually taken to be the Hartree-Fock ground state. We parametrize our trial wavefunction $|\phi\rangle$ as follows:

$$|\phi\rangle = \hat{A} |\Psi\rangle, \hat{A} = A_0 + \sum_{ijkl} A_{jl}^{ik} \hat{c}_i^\dagger \hat{c}_k \hat{c}_j^\dagger \hat{c}_l \quad (6)$$

where $|\Psi\rangle$ is our single-Slater determinant reference wavefunction and \hat{A} is an arbitrary 2-body excitation operator *which need not be expressed in the natural orbital basis of $|\Psi\rangle$* . This latter condition is important for understanding our algorithm, since \hat{A} need not excite only from occupied to virtual orbitals as is often assumed in CISD or CCSD derivations. The constant A_0 is added to \hat{A} to guarantee that the original reference function is contained within our ansatz. Because of the spin-invariance of non-relativistic molecular Hamiltonians, we can constrain our ansatz to preserve the number of α and β electrons. For this reason, our actual ansatz is

$$\hat{A} = A_0 + \hat{A}_{\alpha\alpha} + \hat{A}_{\alpha\beta} + \hat{A}_{\beta\beta}. \quad (7)$$

where the operators $\hat{A}_{\alpha\alpha}$, $\hat{A}_{\alpha\beta}$ and $\hat{A}_{\beta\beta}$ act on two spin-up, one spin-up and one spin-down, or two spin-down electrons, respectively. In what follows, we will ignore this simplification, as it needlessly complicates notation.

Noting that A_{jl}^{ik} is a rank-4 tensor, we apply the THC-DF decomposition from [1] and express it as

$$A_{jl}^{ik} = \sum_{a,b=1}^{P_A} \chi_{ia} \chi_{ka} \mathcal{Z}_{ab} \chi_{jb} \chi_{lb} \quad (8)$$

If we let $P_A = L^2$, we can represent any 2-body excitation operator \hat{A} and hence any CISD wavefunction $|\phi\rangle$. For smaller values of P_A , we do not span the entire CISD space but nonetheless have a very efficient and interesting representation of our trial wavefunction. As we will show in our examples, we are able to obtain excellent approximations to the CISD energy with $P_A < L$.

To obtain an approximate ground state energy, we substitute our trial wavefunction into the standard variational energy expression to obtain

$$E_\Phi = \langle \Phi | \hat{H} | \Phi \rangle / \langle \Phi | \Phi \rangle \quad (9)$$

$$= \left(\langle \hat{A} \hat{H}_1 \hat{A} \rangle + \langle \hat{A} \hat{H}_2 \hat{A} \rangle \right) / \langle \hat{A} \hat{A} \rangle \quad (10)$$

where expectation values are taken over the reference wavefunction $|\Psi\rangle$. The expressions on the right-hand-side of Eq. (10) look formidable because they involve calculating the expectation value of 4- to 6-body operators. It is at this

stage that the THC-DF decomposition becomes crucially important because it allows us to evaluate the quantities in Eq. (10) with only $O(L^4)$ operations. To see how, let us consider the most complicated term in Eq. (10), $\langle \hat{A} \hat{H}_2 \hat{A} \rangle$, which involves the evaluation of a 6-body operator. Plugging in our decompositions in Eqs. (3) and (8), we obtain

$$\langle \hat{A} \hat{H}_2 \hat{A} \rangle = \sum (\chi_{ia} \chi_{ka} \mathcal{Z} \chi_{jb} \chi_{lb}) (x_{mc} x_{oc} Z_{cd} x_{nd} x_{pd}) (\chi_{qe} \chi_{se} \mathcal{Z}_{ef} \chi_{rf} \chi_{tf}) \langle \hat{c}_i^\dagger \hat{c}_k \hat{c}_j^\dagger \hat{c}_l \hat{c}_m^\dagger \hat{c}_o \hat{c}_n^\dagger \hat{c}_p \hat{c}_q^\dagger \hat{c}_s \hat{c}_r^\dagger \hat{c}_t \rangle \quad (11)$$

where we sum over all repeated indices.

Using the canonical fermionic anticommutation relations, we can express the operator in Eq. (11) in canonical ordering, with all creation operators on the left and all annihilation operators on the right. Once the operators are canonically ordered, their expectation values over the single Slater determinant wavefunction $|\Psi\rangle$ can be written exactly as antisymmetrized products of the 1-particle reduced density matrix [19, 20],

$${}^1D_{ik} = \langle \Psi | \hat{c}_i^\dagger \hat{c}_k | \Psi \rangle. \quad (12)$$

For instance, we can write

$$\langle \hat{c}_i^\dagger \hat{c}_j^\dagger \hat{c}_l \hat{c}_k \rangle = {}^1D_k^i {}^1D_l^j - {}^1D_l^i {}^1D_k^j. \quad (13)$$

Expressing the expectation value of an n -body operator over the wavefunction $|\Psi\rangle$ will require $n!$ separate terms, so the result will be lengthy.

The question that immediately arises is whether these lengthy expressions can be evaluated efficiently. At this point, the THC-DF decompositions employed in Eqs. (3) and (8) become extremely important. Because both the Hamiltonian and excitation operators are decomposed in terms of rank-2 tensors, a given index can be connected to *at most* three other indices. For instance, the index a in Eq. (11) is involved only in the terms χ_{ai}, χ_{ak} and \mathcal{Z}_{ab} . As a result, we can hold i, k and b constant and perform the sum over a , storing the result

$$X_{ikb} = \sum_a \chi_{ai} \chi_{ak} \mathcal{Z}_{ab} \quad (14)$$

for future use. When we implement all the terms in Eq. (11), we find that we can evaluate all terms in our expression by summing over a single index while holding at most three others constants. Provided that both P_H and P_A scale as $O(L)$, this fact entails that the expected energy can be evaluated in $O(L^4)$ operations, albeit with a large prefactor. Furthermore, when we take the analytic gradient of our energy with respect to the variational parameters χ_{ai} and \mathcal{Z}_{ab} , we find that this result also requires only $O(L^4)$ operations. An illustrative example of the terms involved in our calculation is provided in the Appendix.

Our final algorithm is implemented as follows: We begin with some initial guess of the variational parameters χ_{ia} and \mathcal{Z}_{ab} . Currently, we have found that selecting a random matrix χ_{ia} and setting $\mathcal{Z} = 0$ appears to work well, especially when P_A is sufficiently large. For small values of $P_A < 5$, the minimization can become trapped in local minima, but sampling a small number of initial conditions guarantees that a repeatable global minimum will be found. We then perform a quasi-Newton minimization on the expected energy E_Φ in Eq. (9), seeking to minimize it with respect to our variational parameters χ_{ia} and \mathcal{Z}_{ab} using the LBFGS optimization algorithm described in [21, 22]. The minimization converged to within 1 mH of its final result within 300 iterations for all the molecules studied and we expect that minimization time can be greatly reduced by using better initial guesses and better approximations to the Hessian in the LBFGS algorithm. Nonetheless, since both the calculation of the energy and its gradient require $O(L^4)$ operations, our minimization algorithm can likewise be accomplished in $O(L^4)$ time. Upon minimization, the energy we obtain will be an upper bound on the CISD energy and will approach the CISD result as the number of auxiliary functions P_A is increased.

Our first set of results in Table I shows the electronic energies obtained for a selection of small molecules in a minimal STO-6G basis set. The RHF and CISD results and the ERIs are calculated using Gaussian09 [23]. The remainder of the table shows the energies calculated by our method as a function of the size of the auxiliary basis P_A . We see that our method converges to the CISD result as we increase the size of the auxiliary basis. More surprising is how few auxiliary functions are needed to obtain a very accurate approximation to the CISD result. For instance, almost 95% of the CISD correlation energy in the molecule LiH could be recovered by our algorithm with only *two* auxiliary functions $P_A = 2$. For other molecules a larger number of auxiliary functions were required to ensure good convergence to the CISD answer. But in all cases, in order to converge recover 98% of the correlation energy present in the exact CISD result, no more than $P_A = 6$ auxiliary functions were required. This result is important because P_A must scale as $O(L)$ to ensure that our overall algorithm retains a scaling of $O(L^4)$. Note that for large P_A , our algorithm recovers slightly *more* correlation energy than is recovered by the CISD algorithm due to spin-contamination in our algorithm.

	E_g (H)	Correlation Energy (mH)				
Molecule	E_{HF}	CISD	$P_A = 2$	$P_A = 4$	$P_A = 6$	$P_A = 10$
BH	-27.1484	55.7	37.8	55.4	55.8	55.9
LiH	-8.9182	20.9	20.1	21.0	21.0	21.1
BeH2	-19.0803	34.5	29.7	32.2	34.7	34.8
CH2	-44.7828	58.1	35.7	55.0	57.2	58.2
HF	-103.1456	66.7	65.0	66.1	66.4	66.7
H2O	-84.7683	50.7	46.0	47.2	50.0	50.5

TABLE I: Ground state energies for a variety of small molecules in the STO-6G basis set. As the number of auxiliary functions, P_A , is increased, the energy obtained by our algorithm approach the exact CISD result. In all cases, our method with $P_A = 6$ recovers more than 98% of the CISD correlation energy.

	E_g (H)	Correlation Energy (mH)				
Molecule	E_{HF}	CISD	$P_A = 2$	$P_A = 4$	$P_A = 6$	$P_A = 10$
BH	-27.2559	59.9	32.1	52.8	59.0	59.8
LiH	-8.9475	19.0	17.5	18.6	19.1	19.2
BeH2	-19.1161	39.4	16.8	34.4	38.8	39.3
CH2	-44.8861	84.3	36.6	58.2	77.2	81.6
HF	-103.6346	144.9	67.8	128.1	139.6	143.3
H2O	-85.0717	130.2	47.3	111.0	121.0	125.8

TABLE II: Ground state energies for a variety of small molecules in the 6-31G basis set. As the number of auxiliary functions, P_A , is increased, the energy obtained by our algorithm approach the exact CISD result. In all cases, our method with $P_A = 10$ recovers more than 96% of the CISD correlation energy.

While the CISD algorithm is spin-projected such that only singlet states are included, our algorithm allows mixing into states of other spin-multiplicities, yielding energies that are ever so slightly lower than the spin-pure result. In all cases, this effect was negligible, amounting to less than 0.3 mH.

Table II shows the same results for the same molecules using the larger 6-31G basis. Again, we see a rapid convergence to the exact CISD result as a function of P_A . These results also support the inference that only $P_A = O(L)$ auxiliary functions are needed to converge to the exact result, since in all cases studied, setting $P_A = L$ recovers $> 97\%$ of the exact CISD correlation energy.

Finally, Figure 1 offers further support for our contention that our algorithm scales as $O(L^4)$. This plot shows the computational cost of calculating the energy and gradient of a system of N non-interacting H_2 molecules as a function of N . Because the molecules are strictly non-interacting, $P'_H = NP_H$, where P'_H and P_H are the number of auxiliary functions needed to describe the composite and single system, respectively. Similarly, we use $P'_A = NP_A$ auxiliary functions to represent our excitation operator. As a result, this system provides the perfect test case for measuring the computational scaling of our algorithm. The result in Fig. 1 shows that our algorithm scales as $O(L^4)$. This measurement is consistent with an examination of our code. The slow step in our algorithm is a series of nested do-loops that sum four indices over the range $1 \dots P_A$ or $1 \dots P_H$. Provided that P_A and P_H both scale as $O(L)$, our algorithm then has a complexity of $O(L^4)$, as we observe.

There are numerous areas of future research. One of the most important issues is performing the efficient THC-DF decomposition of the Hamiltonian. Although the authors of [1] introduced a method to decompose a molecular Hamiltonian in terms of $P_H = O(L)$ auxiliary functions using $O(L^4)$ operations, they recognized that there were very large prefactors involved in this process. For this reason, it is worth investigating whether we can construct such a decomposition more efficiently, even if the construction process retains the same asymptotic scaling.

A related issue is the prefactor associated with the electronic structure algorithm presented in this paper. The slow step of our algorithm is the calculation of the energy and gradient of the 6-body term $\langle \hat{A} \hat{H}_2 \hat{A} \rangle$. Since this expectation value involves six creation and six annihilation operators, there are $6! = 720$ possible matchings of the operators, each of which contributes a unique term to the energy and gradient. The code required to perform such calculations is immense; our final program contained over 500,000 lines of FORTRAN code even when effort was made to store intermediate results and eliminate redundancy. Major improvement could be made to our algorithm either by developing analytic methods to reduce the cost of energy/gradient evaluation or by developing methods which




FIG. 1: The computational cost of each energy/gradient calculation as a function of the number of basis function L for a system of N noninteracting H_2 molecules in the STO-6G basis. The blue Xs show the numerical results; for comparison, the dashed line shows L^4 scaling. As expected, our algorithm scales asymptotically as $O(L^4)$

identify and approximate negligible terms in our equations.

A final major area of research is the extension of our method to coupled-cluster algorithms. One of the key elements of coupled cluster theory is that the excitation operator T_2 must only connect occupied orbitals to virtual orbitals; there can be no non-zero elements which act within the occupied or virtual subspaces or which connect virtual orbitals to occupied orbitals. However, if we make this assumption, it is not clear that the excitation operator T_2 can still be efficiently written in the THC-DF form, which is vital to the efficiency of our algorithm. In other words, the two approximations which are crucial to an THC-DF-based coupled-cluster theory may be mutually incompatible. Future studies will have to assess whether this is indeed the case or whether either theory can be modified to produce an efficient coupled-cluster algorithm. It is also interesting that it has already been shown in a different context that generalized coupled-cluster methods, which do not make the occupied/unoccupied distinction, can exactly represent the ground state many-body wavefunction[24]. Our results provide additional motivation for investigating generalized coupled-cluster methods.

In conclusion, we have described an $O(L^4)$ electronic structure algorithm which can be tuned in its accuracy through the value of the parameter P_A . When this parameter is 0, we recover the Hartree-Fock result. When the parameter is $O(L)$, we rapidly approach the exact CISD result. Although the percentage of correlation energy recovered by the CISD approach decreases to 0 as a function of the system size due to its lack of size extensivity, the many known approaches to recovering size-extensivity within CISD should be applicable to our method, thereby resotring its utility even for large systems[18]. Our hope is that the efficiency of this algorithm will make it computationally competitive with existing low-scaling methods like Hartree-Fock and density functional theory, bringing explicitly correlated wavefunction-based methods into the realm of practical calculation even for large systems.

Acknowledgments

NS and WY would like to acknowledge support from the National Science Foundation (CHE-09-11119). HvA thanks the FWO-Flanders and Duke University for support.

I. APPENDIX

The total variational energy is given by

$$E_\Phi = \left(\langle \hat{A} \hat{H}_1 \hat{A} \rangle + \langle \hat{A} \hat{H}_2 \hat{A} \rangle \right) / \langle \hat{A} \hat{A} \rangle \quad (15)$$

where the \hat{A} operator is given by

$$\hat{A} = A_0 + \hat{A}_{\alpha\alpha} + \hat{A}_{\alpha\beta} + \hat{A}_{\beta\beta} \quad (16)$$

and where the operators \hat{H}_1 and \hat{H}_2 conserve spin such that they can be written as

$$\hat{H}_1 = \hat{H}_1^\alpha + \hat{H}_1^\beta \quad (17)$$

$$\hat{H}_2 = \hat{H}_2^{\alpha\alpha} + \hat{H}_2^{\alpha\beta} + \hat{H}_2^{\beta\beta} \quad (18)$$

For our demonstration, we will consider only the contributions to the total energy that come from the term

$$E = \langle A_0 \hat{H}_1^\alpha \hat{A}_{\alpha\alpha} \rangle \quad (19)$$

Using the THC-DF decomposition for the excitation operator, we obtain the expression

$$E = A_0 \sum_{ikabjlmn} h_{ik} \chi_{aj} \chi_{al} \mathcal{Z}_{ab} \chi_{bm} \chi_{bn} \langle \hat{c}_i^\dagger \hat{c}_k \hat{c}_j^\dagger \hat{c}_l \hat{c}_m^\dagger \hat{c}_n \rangle \quad (20)$$

To evaluate the expectation value in Eq. (20), we use the canonical fermionic anticommutation relations. For convenience, we define the 1-hole reduced density matrix ${}^1Q_k^i = \delta_k^i - {}^1D_k^i$. After some tedious algebra, we obtain

$$\begin{aligned} \langle \hat{c}_i^\dagger \hat{c}_k \hat{c}_j^\dagger \hat{c}_l \hat{c}_m^\dagger \hat{c}_n \rangle &= {}^1D_k^i {}^1D_l^j {}^1D_n^m + {}^1D_k^i {}^1D_n^j {}^1Q_m^l + {}^1D_l^i {}^1Q_j^k {}^1D_n^m \\ &\quad + {}^1D_n^i {}^1Q_m^k {}^1D_l^j - {}^1D_l^i {}^1Q_m^k {}^1D_n^j + {}^1D_n^i {}^1Q_j^k {}^1Q_m^l \end{aligned} \quad (21)$$

If we plug Eq. (21) into Eq. (20), we obtain

$$\begin{aligned} E &= A_0 \sum_{ab} \left(\text{Tr}[h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{aa} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{bb} \right. \\ &\quad + \text{Tr}[h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{ab} \mathcal{Z}_{ab} [\chi \cdot Q \cdot \chi^\dagger]_{bb} \\ &\quad + [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{aa} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{bb} \\ &\quad + [\chi \cdot D \cdot \chi^\dagger]_{aa} \mathcal{Z}_{ab} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{bb} \\ &\quad - [\chi \cdot D \cdot h \cdot Q \cdot \chi^\dagger]_{ab} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{ab} \\ &\quad \left. + [\chi \cdot Q \cdot \chi^\dagger]_{ab} \mathcal{Z}_{ab} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{ab} \right) \end{aligned} \quad (22)$$

The operations in Eq. (20) consist of matrix multiplications and summations over the indices a and b . The matrices h , D , and Q are $L \times L$ matrices while the matrix χ is a $P_A \times L$ matrix. Provided that $P_A = O(L)$, all of the matrix multiplications will require $O(L^3)$ operations. Finally, the sum over indices a and b runs from 1 to P_A . As long as P_A scales linearly with L the sum over a and b can be performed in $O(L^2)$ operations.

The value of the gradient can be calculated by differentiating Eq. (22) with respect to the various parameters.

$$dE/dA_0 = E/A_0 \quad (23)$$

$$dE/d\chi_{ai} = A_0 \sum_b \left(\text{Tr} [h \cdot D] [D \cdot \chi^\dagger]_{ia} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{bb} \right. \quad (24)$$

$$\begin{aligned} & + \text{Tr} [h \cdot D] [\chi \cdot D]_{ai} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{bb} \\ & + \text{Tr} [h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{bb} \mathcal{Z}_{ba} [D \cdot \chi^\dagger]_{ia} \\ & + \text{Tr} [h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{bb} \mathcal{Z}_{ba} [\chi \cdot D]_{ai} \\ & + \text{Tr} [h \cdot D] [D \cdot \chi^\dagger]_{ib} \mathcal{Z}_{ab} [\chi \cdot Q \cdot \chi^\dagger]_{bb} \\ & + \text{Tr} [h \cdot D] [\chi \cdot D]_{bi} \mathcal{Z}_{ba} [\chi \cdot Q \cdot \chi^\dagger]_{aa} \\ & + \text{Tr} [h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{ba} \mathcal{Z}_{ba} [Q \cdot \chi^\dagger]_{ia} \\ & + \text{Tr} [h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{ba} \mathcal{Z}_{ba} [\chi \cdot Q]_{ai} \\ & + [Q \cdot h \cdot D \cdot \chi^\dagger]_{ia} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{bb} \\ & + [\chi \cdot Q \cdot h \cdot D]_{ai} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{bb} \\ & + [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{bb} \mathcal{Z}_{ba} [D \cdot \chi^\dagger]_{ia} \\ & + [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{bb} \mathcal{Z}_{ba} [\chi \cdot D]_{ai} \\ & + [D \cdot \chi^\dagger]_{ia} \mathcal{Z}_{ab} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{bb} \\ & + [\chi \cdot D]_{ai} \mathcal{Z}_{ab} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{bb} \\ & + [\chi \cdot D \cdot \chi^\dagger]_{bb} \mathcal{Z}_{ba} [Q \cdot h \cdot D \cdot \chi^\dagger]_{ia} \\ & + [\chi \cdot D \cdot \chi^\dagger]_{bb} \mathcal{Z}_{ba} [\chi \cdot Q \cdot h \cdot D]_{ai} \\ & - [D \cdot h \cdot Q \cdot \chi^\dagger]_{ib} \mathcal{Z}_{ab} [\chi \cdot D \cdot \chi^\dagger]_{ab} \\ & - [\chi \cdot D \cdot h \cdot Q]_{bi} \mathcal{Z}_{ba} [\chi \cdot D \cdot \chi^\dagger]_{ba} \\ & - [\chi \cdot D \cdot h \cdot Q \cdot \chi^\dagger]_{ab} \mathcal{Z}_{ab} [D \cdot \chi^\dagger]_{ib} \\ & - [\chi \cdot D \cdot h \cdot Q \cdot \chi^\dagger]_{ba} \mathcal{Z}_{ba} [\chi \cdot D]_{bi} \\ & + [Q \cdot \chi^\dagger]_{ib} \mathcal{Z}_{ab} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{ab} \\ & + [\chi \cdot Q]_{bi} \mathcal{Z}_{ba} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{ba} \\ & + [\chi \cdot Q \cdot \chi^\dagger]_{ab} \mathcal{Z}_{ab} [Q \cdot h \cdot D \cdot \chi^\dagger]_{ib} \\ & + [\chi \cdot Q \cdot \chi^\dagger]_{ba} \mathcal{Z}_{ba} [\chi \cdot Q \cdot h \cdot D]_{bi} \end{aligned}$$

$$dE/d\mathcal{Z}_{ab} = A_0 \sum_{ab} \left(\text{Tr} [h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{aa} [\chi \cdot D \cdot \chi^\dagger]_{bb} \right. \quad (25)$$

$$\begin{aligned} & + \text{Tr} [h \cdot D] [\chi \cdot D \cdot \chi^\dagger]_{ab} [\chi \cdot Q \cdot \chi^\dagger]_{bb} \\ & + [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{aa} [\chi \cdot D \cdot \chi^\dagger]_{bb} \\ & + [\chi \cdot D \cdot \chi^\dagger]_{aa} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{bb} \\ & - [\chi \cdot D \cdot h \cdot Q \cdot \chi^\dagger]_{ab} [\chi \cdot D \cdot \chi^\dagger]_{ab} \\ & + [\chi \cdot Q \cdot \chi^\dagger]_{ab} [\chi \cdot Q \cdot h \cdot D \cdot \chi^\dagger]_{ab} \end{aligned}$$

Notice that many of the quantities involved in the above equations are repeated. These can be calculated and stored for later use, to speed computation at the expense of a larger memory requirement. The expressions above provide a representative example of the kinds of calculations involved in our computation.

- [2] I. Levine, *Quantum Chemistry* (Prentice Hall, Upper Saddle River, 2000).
- [3] A. Szabo and N. Ostlund, *Modern Quantum Chemistry* (Dover, Mineola, 1996).
- [4] R. Bartlett, Ann. Rev. Phys. Chem. **32**, 359 (1981).
- [5] M. Head-Gordon, J. Phys. Chem. **100**, 13213 (1996).
- [6] R. Bartlett and M. Musial, Rev. Mod. Phys. **79**, 291 (2007).
- [7] D. Mazziotti, Phys. Rev. A **57**, 4219 (1998).
- [8] D. Mazziotti, J. Chem. Phys. **121**, 10957 (2004).
- [9] P.-A. Malmqvist and B. O. Roos, Chem. Phys. Lett. **155**, 189 (1989).
- [10] A. Cohen, P. Mori-Sánchez, and W. Yang, Chem. Rev. **112**, 289 (2012).
- [11] O. Vahtras, J. Almlöf, and M. Feyereisen, Chem. Phys. Lett. **213**, 514 (1993).
- [12] K. Eichkorn, O. Treutler, H. Öhm, M. Häser, and R. Ahlrichs, Chem. Phys. Lett. **264**, 573 (1997).
- [13] R. Kendall and H. Früchtel, Theor. Chem. Acc. **97**, 158 (1997).
- [14] F. Weingend and M. Haser, Theor. Chem. Acc. **97**, 331 (1997).
- [15] T. Martinez and E. Carter, J. Chem. Phys. **98**, 7081 (1993).
- [16] T. Martinez and E. Carter, J. Chem. Phys. **100**, 3631 (1994).
- [17] T. Martinez and E. Carter, J. Chem. Phys. **102**, 7564 (1995).
- [18] W. Duch and G. Dierksen, J. Chem. Phys. **101**, 3018 (1994).
- [19] D. Mazziotti, Int. J. Quantum Chem. **70**, 557 (1998).
- [20] M. Benayoun, A. Lu, and D. Mazziotti, Chem. Phys. Lett. **387**, 485 (2004).
- [21] D. Liu and J. Nocedal, Mathematical Programming B **45**, 503 (1989).
- [22] C. Zhu, R. Byrd, and J. Nocedal, ACM Transactions on Mathematical Software **23**, 550 (1997).
- [23] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, B. Mennucci, G. A. Petersson, et al., *Gaussian 09 Revision A.1*, gaussian Inc. Wallingford CT 2009.
- [24] M. Nooijen, Phys. Rev. Lett. **84**, 2108 (2000).